arXiv:1607.07028v1 [stat.ME] 24 Jul 2016

A Statistical Model for the Analysis of Beta Values in DNA Methylation Studies

Leonie Weinhold[1], Simone Wahl[2], Matthias Schmid[1]

[1]Department of Medical Biometry, Informatics and Epidemiology
University of Bonn, Sigmund-Freud-Str. 25, D-53127 Bonn, Germany
Email: weinhold@imbie.uni-bonn.de

[2]Research Unit of Molecular Epidemiology, Helmholtz Zentrum München
Ingolstädter Landstr. 1, D-85764 Neuherberg, Germany

Abstract

Background: The analysis of DNA methylation is a key component in the development of personalized treatment approaches. A common way to measure DNA methylation is the calculation of beta values, which are bounded variables of the form $M/(M+U)$ that are generated by Illumina's 450k BeadChip array. The statistical analysis of beta values is considered to be challenging, as traditional methods for the analysis of bounded variables, such as M-value regression and beta regression, are based on regularity assumptions that are often too strong to adequately describe the distribution of beta values.

Results: We develop a statistical model for the analysis of beta values that is derived from a bivariate gamma distribution for the signal intensities M and U. By allowing for possible correlations between M and U, the proposed model explicitly takes into account the data-generating process underlying the calculation of beta values.

Conclusion: The proposed model can be used to improve the identification of associations between beta values and covariates such as clinical variables and lifestyle factors in epigenome-wide association studies. It is as easy to apply to a sample of beta values as beta regression and M-value regression.

1 Background

The analysis of DNA methylation has become of considerable interest in biomedical research, as epigenetic studies have shown numerous associations between methylation levels and diseases such as cancer and cardiovascular disease [2, 27, 24, 30, 32]. Today, most research focuses on the cytosine-guanine dinucleotide ("CpG") sites of the DNA, which are the locations where methylation is primarily found in humans [20]. One of the most widely used techniques to measure DNA methylation is the Illumina Infinium HumanMethylation450 BeadChip array, which covers approximately 450,000 CpG sites. At each CpG site, methylation is quantified by the beta value $b := M/(M+U+a)$, where $M > 0$ and $U > 0$ denote the methylated and unmethylated signal intensities, respectively, measured by the Illumina 450k array. The offset $a \geq 0$ is usually set equal to 100 and is added to $M+U$ to stabilize beta values when both M and U are small.

An important goal of methylation analysis is to identify DNA regions where methylation is associated with disease status, lifestyle factors and other clinical or sociodemographic variables [5, 31, 7, 28]. This is often achieved by fitting site-wise regression models with dependent variable b and a vector of covariates $\boldsymbol{X}$ that may also include potential confounders. After model fitting, a common strategy is to carry out downstream hypothesis tests to identify those CpG sites that show significant associations between methylation status and the variables of interest.

Because, by definition, b is bounded between 0 and 1, Gaussian regression with untransformed beta values is problematic in the context of DNA methylation analysis. In particular, the variance of b is usually smaller near the boundaries than near the middle of the interval (0,1), implying that the homoscedasticity assumption in Gaussian regression is violated [9, 15, 26]. To address this problem, several modeling strategies have been developed, including Gaussian regression with logit-transformed

beta values ("M-values", [9]) and generalized regression models for untransformed bounded responses, e.g. beta regression [10]. Regarding the analysis of DNA methylation, both strategies are intrinsically problematic: In case of M-value regression, the assumptions of a Gaussian model are often not met despite the transformation of the data, and the interpretation of the coefficient estimates is only possible on the transformed scale but not on the original scale of b [10, 3]. Beta regression, on the other hand, requires the ratio $M/(M+U+a)$ to follow a beta distribution, implicitly assuming that the variables M and U are independently gamma distributed [6]. While M and U can indeed be described by gamma distributed random variables [29, 23], the independence assumption for the two signal intensities is often not met in practice. For example, Laird [15] reported that the methylated and unmethylated signal intensities, as produced by the Illumina 450k array, are usually positively correlated. These issues, along with the results of two recent empirical studies [23, 31], suggest that both M-value regression and beta regression need to be improved to describe the distribution of b in a statistically sound way.

To address this problem, we propose a novel analysis technique for beta values that relaxes the independence assumption between the signal intensities M and U. The idea is to start with a model for the bivariate distribution of M and U and to derive the probability density function of the ratio $M/(M+U)$. This function is subsequently used to construct the log-likelihood function of a generalized regression model that relates beta values to linear functions of the covariates. Because estimation of the model parameters is based on the maximum likelihood principle, asymptotic confidence intervals and normally distributed test statistics can be derived by evaluating the inverse of the observed information matrix. This strategy allows for downstream hypothesis tests on the associations between a covariate of interest and the methylation status at individual CpG sites. For the rest of this paper, we will refer to the proposed model as "RCG" (*R*atio of *C*orrelated *G*ammas) model.

2 Methods and Results

In Section 2.1 we introduce basic notation and definitions. Section 2.2 briefly reviews beta regression and M-value regression and discusses the limitations of the two methods. In Section 2.3 the proposed RCG model for the analysis of beta values is derived. Section 2.4 provides details on model fitting and on the construction of downstream hypothesis tests.

2.1 Notation and Definitions

At each CpG site, the Illumina 450k array produces a sample of methylated and unmethylated signal intensities $(M_i, U_i)_{i=1,\ldots,n}$, where n is the number of analyzed persons. The corresponding set of beta values is calculated by $b_i = M_i/(M_i + U_i + a)$, $i = 1, \ldots, n$. To facilitate the derivation of distributional results, we will set $a = 0$ throughout this section. The predictor variable(s) of interest and the confounding variables are collected in vectors $\boldsymbol{X}_i = (1, X_{i1}, \ldots, X_{ip})^\top$, $i = 1, \ldots, n$. For each CpG site, the aim is to analyze the associations between the variables in $\boldsymbol{X}$ and the methylation status b.

Following [29] and [23], we assume that the stochastic behavior of the signal intensities M and U can be described by gamma distributed random variables with densities

$$f_M(m) = \frac{\lambda_m}{\Gamma(\alpha_m)} (\lambda_m m)^{\alpha_m - 1} \exp(-\lambda_m m)\,, \tag{1}$$

$$f_U(u) = \frac{\lambda_u}{\Gamma(\alpha_u)} (\lambda_u u)^{\alpha_u - 1} \exp(-\lambda_u u)\,, \tag{2}$$

where α_m, α_u and λ_m, λ_u are the shape and rate parameters of f_M and f_U, respectively. From (1) it follows that the means and variances of M, U are given by α_m/λ_m, α_u/λ_u and α_m/λ_m^2, α_u/λ_u^2, respectively [1].

2.2 Regression Models for the Analysis of Beta Values

Since the ratio $b = M/(M+U)$ is bounded between 0 and 1, it has been argued that a linear regression model of the form

$$b = \boldsymbol{X}^\top \gamma + \epsilon\,, \quad \gamma \in \mathbb{R}^{p+1}\,, \quad \epsilon \sim \mathcal{N}(0, \sigma^2)\,, \tag{3}$$

is not appropriate to model DNA methylation. In particular, the variance of b is usually smaller near the boundaries than near the middle of the interval (0,1), implying that the homoscedasticity assumption $\mathrm{var}(\epsilon) = \sigma^2$ is violated [9].

In view of this problem, several statistical models for bounded response variables have been developed (see [26] for an overview). A simple approach is to calculate logit-transformed beta values ("M-values", [9]) and to fit a linear regression model of the form

$$\log_2\left(\frac{b}{1-b}\right) = \boldsymbol{X}^\top \gamma + \epsilon\,, \quad \epsilon \sim \mathcal{N}(0, \sigma^2)\,. \tag{4}$$

Although this strategy has become popular in the analysis of DNA methylation, it has the drawback that the methylation status (as quantified by the value of b) is not analyzed on its original scale but on a transformed scale [10]. Furthermore, as shown by Wahl *et al.* [31], the empirical distribution of logit-transformed beta values usually deviates from normality.

An alternative approach that operates on the untransformed scale of b is *beta regression*, which is characterized by a beta distributed outcome variable with probability density function

$$\varphi(b) = \frac{\Gamma(\phi)}{\Gamma(\mu\phi)\Gamma((1-\mu)\phi)}\, b^{\mu\phi-1}\,(1-b)^{(1-\mu)\phi-1}\,, \tag{5}$$

where μ and ϕ denote the mean and precision parameters, respectively, of the probability density function φ. The predictor-response relationship is usually defined by a monotone increasing link function $g(\cdot)$ and by the model equation $g(\mu|\boldsymbol{X}) = \boldsymbol{X}^\top\gamma$ [10]. A common choice for g is the logit transformation $\log(\mu/(1-\mu))$. Since the variance of a beta distributed random variable is given by $\mu(1-\mu)/(1+\phi)$, beta regression accounts for heteroscedasticity and for small variances near the boundaries of the interval (0,1). On the other hand, a major shortcoming of (5) in the context of DNA methylation analysis is that the signal intensities M and U are implicitly assumed to be independent and to share a common rate parameter. Under these assumptions, the ratio $b = M/(M+U)$ can be shown to follow a beta distribution ([6], Chapter 9). The independence assumption, however, cannot be confirmed by empirical findings, which show that the signal intensities obtained from the Illumina 450k array are often positively correlated (see [15]).

2.3 A Statistical Model for the Ratio of Correlated Gamma Distributed Random Variables

To overcome the problems associated with M-value regression and beta regression, we propose a statistical model ("Ratio of Correlated Gammas (RCG) model") that is based on the bivariate distribution of the signal intensities M and U. In contrast to beta regression, we assume that M and U are not independent but can be described by a bivariate gamma distribution with probability density function

$$\begin{aligned} f_{M,U}(m,u) \;=\; & \frac{(\lambda_m\lambda_u)^\alpha}{(1-\rho)\,\Gamma(\alpha)}\left(\frac{mu}{\rho\,\lambda_m\lambda_u}\right)^{\frac{\alpha-1}{2}} \exp\left(-\frac{\lambda_m m}{1-\rho}\right) \\ & \times \exp\left(-\frac{\lambda_u u}{1-\rho}\right) I_{\alpha-1}\left(\frac{2\sqrt{\rho\lambda_m\lambda_u mu}}{1-\rho}\right)\,, \end{aligned} \tag{6}$$

where $\lambda_m, \lambda_u, \alpha > 0$, $0 < \rho < 1$, and $I_{\alpha-1}$ is the modified Bessel function of the first kind of order $\alpha-1$. The distribution in (6) is due to Kibble [14] and is often referred to as "Wicksell-Kibble bivariate gamma distribution" [1]. As stated in various articles and monographs (e.g. [16]), the marginal

densities f_M, f_U of M and U, respectively, are given by

$$f_M(m) = \frac{\lambda_m}{\Gamma(\alpha)} (\lambda_m m)^{\alpha-1} \exp(-\lambda_m m), \tag{7}$$

$$f_U(u) = \frac{\lambda_u}{\Gamma(\alpha)} (\lambda_u u)^{\alpha-1} \exp(-\lambda_u u). \tag{8}$$

The equations in (7) and (8) imply that M and U are gamma distributed random variables with a common shape parameter α and with means and variances given by α/λ_m, α/λ_u and α/λ_m^2, α/λ_u^2, respectively. The restriction to a common shape parameter ensures that all measured signal intensities refer to probability density functions sharing the same basic form. On the other hand, the unequal rate parameters λ_m and λ_u guarantee sufficient flexibility in modeling the differences in the marginal densities of M and U (see (11) and (12)). It can further be shown that the Pearson correlation of M and U is equal to ρ, implying that (6) imposes a correlation structure on the two signal intensities (see [1]).

In the next step, the distribution of the ratio $b = M/(M+U)$ is derived:

Proposition 1. Let the distribution of (M, U) be defined by the probability density function in (6). Then the ratio $b = M/(M+U)$ follows a univariate distribution with probability density function

$$\begin{aligned} f_b(b) &= \frac{\Gamma(2\alpha)}{\Gamma^2(\alpha)} (\lambda_m \lambda_u)^{\alpha} (1-\rho)^{\alpha} \, (b(1-b))^{\alpha-1} \\ &\quad \times \frac{(\lambda_m b + \lambda_u (1-b))}{\Big((\lambda_m b + \lambda_u (1-b))^2 - 4\rho\lambda_m\lambda_u b(1-b) \Big)^{\alpha+0.5}} \, . \end{aligned} \tag{9}$$

Proof: The proof of Proposition 1, which is related to the work of Nadarajah and Kotz [18], is given in the appendix.

The result stated in Proposition 1 can be used to derive the log-likelihood function of a sample of beta values $b_1, \ldots, b_n$:

Proposition 2. For independent sample values $b_1, \ldots, b_n$, the log-likelihood function derived from (9) is given by

$$\begin{aligned} \sum_{i=1}^{n} \log(f_b(b_i; \alpha, \rho, \theta)) &= \sum_{i=1}^{n} \Big[\log(\Gamma(2\alpha)) - 2\log(\Gamma(\alpha)) + \alpha \log(\theta) + \alpha \log(1-\rho) \\ &\quad + \log\left((\theta-1)b_i + 1\right) + (\alpha - 1)\log(b_i(1-b_i)) \\ &\quad - \big(\alpha + 0.5\big) \log\left(\left((\theta-1)b_i + 1\right)^2 - 4\,\rho\,\theta\, b_i(1-b_i) \right) \Big], \end{aligned} \tag{10}$$

where $\theta := \lambda_m/\lambda_u$.

Proof: See appendix.

Proposition 2 implies that the log-likelihood function derived from (9) is a function of the mean ratio $\theta = \lambda_m/\lambda_u = \mathrm{E}(U)/\mathrm{E}(M)$.

To quantify the associations between the covariates $\boldsymbol{X}$ and the signal intensities M and U, we consider linear predictors $\boldsymbol{X}^\top \zeta_m$ and $\boldsymbol{X}^\top \zeta_u$, $\zeta_m, \zeta_u \in \mathbb{R}^{p+1}$, that relate the vector $\boldsymbol{X} = (1, X_1, \ldots, X_p)^\top$ to the marginal means α/λ_m and α/λ_u, respectively. A convenient link function that guarantees the positivity of λ_m and λ_u is the logarithmic transformation, resulting in the predictor-response relationships

$$\log(\mathrm{E}(M|\boldsymbol{X})) = \log(\alpha) - \boldsymbol{X}^\top \zeta_m, \tag{11}$$

$$\log(\mathrm{E}(U|\boldsymbol{X})) = \log(\alpha) - \boldsymbol{X}^\top \zeta_u, \tag{12}$$

with $\log(\lambda_m) = \boldsymbol{X}^\top \zeta_m$ and $\log(\lambda_u) = \boldsymbol{X}^\top \zeta_u$. Note that the term $\log(\alpha)$ can be incorporated into the intercept terms of the coefficient vectors $\zeta_m = (\zeta_{0m}, \zeta_{1m}, \ldots, \zeta_{pm})^\top$ and $\zeta_u = (\zeta_{0u}, \zeta_{1u}, \ldots, \zeta_{pu})^\top$. The model equations in (11) and (12) are therefore in line with traditional univariate gamma regression approaches that relate the log-transformed mean of the response variable to a linear function of the predictors.

Defining $\gamma = (\gamma_0, \gamma_1, \ldots, \gamma_p)^\top := \zeta_m - \zeta_u$, the mean ratio $\mathrm{E}(U|\boldsymbol{X})/\mathrm{E}(M|\boldsymbol{X})$ can be written as $\theta|\boldsymbol{X} = \exp(\boldsymbol{X}^\top\gamma)$, and the log-likelihood function of a sample $(b_1, \boldsymbol{X}_1^\top), \ldots, (b_n, \boldsymbol{X}_n^\top)$ becomes

$$\begin{aligned}\sum_{i=1}^{n} \log(f_b(b_i, \boldsymbol{X}_i; \alpha, \rho, \gamma)) \;&=\; \sum_{i=1}^{n} \Bigg[\log(\Gamma(2\alpha)) - 2\log(\Gamma(\alpha)) + \alpha\, \boldsymbol{X}_i^T\gamma \\ &+ \alpha\log(1-\rho) + \log\big((\exp(\boldsymbol{X}_i^T\gamma) - 1)\, b_i + 1\big) + (\alpha - 1)\, \log(b_i(1-b_i)) \\ &- (\alpha + 0.5) \log\Big(\big((\exp(\boldsymbol{X}_i^T\gamma) - 1)\, b_i + 1\big)^2 - 4\,\rho \exp(\boldsymbol{X}_i^T\gamma)\, b_i(1-b_i)\Big)\Bigg]. \end{aligned} \tag{13}$$

Equations (11) to (13) define a statistical model in which the association between the methylation status b and the covariates $\boldsymbol{X}$ is quantified by the coefficient vector γ. If $\gamma_k = 0$, $k \in \{1, \ldots, p\}$, the predictor-response relationships in (11) and (12) imply that $\zeta_{km} = \zeta_{ku}$ and $\mathrm{E}(M|\boldsymbol{X}) = \mathrm{E}(U|\boldsymbol{X})$ (provided that the values of the other covariates remain constant). Hence, if $\gamma_k = 0$, the k-th covariate X_k has the same effect on both M and U, implying that X_k is not associated with the methylation status at the CpG site under consideration. On the other hand, large values of $|\gamma_k|$ result from large differences in the coefficients ζ_{km} and ζ_{ku}, implying that DNA methylation varies greatly with the value of X_k. Assessing the hypotheses "$H_0 : \gamma_k = 0$ vs. $H_1 : \gamma_k \neq 0$" is therefore equivalent to a statistical test on the association between b and X_k.

2.4 Estimation and Hypothesis Tests

To obtain a consistent estimator of the coefficient vector γ, the log-likelihood function in (13) needs to be maximized over both γ and the hyperparameters α and ρ. To this purpose, we propose the application of a gradient boosting algorithm with linear base-learning functions, as described in [4]. For given data $(b_i, \boldsymbol{X}_i^\top)_{i=1,\ldots,n}$, gradient boosting is a generic optimizer that minimizes a risk function $\mathcal{R}(f, (b_i, \boldsymbol{X}_i^\top)_{i=1,\ldots,n})$ over an unknown prediction function $f(\boldsymbol{X})$, with the only requirement being the existence of the derivative $\partial\mathcal{R}/\partial f$ [13]. Because the base-learning functions are chosen to be linear in $\boldsymbol{X}$, the space of the prediction function f is restricted to the subspace defined by $f(\boldsymbol{X}) = \boldsymbol{X}^\top\gamma$, implying that estimation of f reduces to the estimation of the coefficient vector γ (see [11] for a detailed description of the algorithm). Furthermore, gradient boosting allows for the additional estimation of the hyperparameters α and ρ [25]. Maximum likelihood (ML) estimates of γ, α and ρ can therefore be obtained by setting $\mathcal{R}$ equal to the negative of the log-likelihood in (13) and by running gradient boosting until convergence.

By standard maximum likelihood arguments, the hypotheses "$H_0 : \gamma_k = 0$ vs. $H_1 : \gamma_k \neq 0$" can be investigated by plugging the ML estimates $\hat{\gamma}$, $\hat{\alpha}$ and $\hat{\rho}$ in the observed information matrix $J(\alpha, \rho, \gamma) = -\sum_{i=1}^{n} \partial^2 \log(f_b(b_i, \boldsymbol{X}_i; \alpha, \rho, \gamma))/\partial^2\gamma$ and by calculating the test statistic

$$Z_k = \hat{\gamma}_k \,/\, \sqrt{J_{kk}^{-1}(\hat{\alpha}, \hat{\rho}, \hat{\gamma})}\,, \quad k \in \{1, \ldots, p\}\,, \tag{14}$$

where J_{kk}^{-1} denotes the k-th diagonal element of J^{-1}. Under the null hypothesis, Z_k is asymptotically standard normally distributed as $n \to \infty$. Details on the calculation of J are given in the appendix.

3 Discussion

The development of statistical models to analyze DNA methylation is the subject of intense and ongoing research [19, 33, 7, 8]. In this paper, we proposed a likelihood-based approach to analyze and infer the associations between covariates and methylation levels in Illumina 450k data. In contrast

to beta regression, the proposed RCG model accounts for possible correlations between methylated and unmethylated signal intensities, thereby increasing the flexibility of the model in describing the distribution of methylation levels at individual CpG sites.

The use of a gradient boosting algorithm to optimize the parameters of the RCG model lays the ground for a variety of additional modeling options. For example, it is straightforward to account for nonlinear covariate effects and to extend the linear predictor in (13) by a set of spline functions. Furthermore, it is possible to embed the RCG model in the GAMLSS framework [22] and to increase its flexibility by relating the parameters α and ρ to separate linear or additive predictors. For details, see [17] and [12].

References

[1] N. Balakrishnan and C.-D. Lai. *Continuous Bivariate Distributions*. Springer, New York, 2 edition, 2009.

[2] A. Bird. DNA methylation patterns and epigenetic memory. *Genes & Development*, 16:6–21, 2002.

[3] C. Bock. Analysing and interpreting DNA methylation data. *Nature Reviews Genetics*, 13: 705–719, 2012.

[4] P. Bühlmann and T. Hothorn. Boosting algorithms: Regularization, prediction and model fitting. *Statistical Science*, 22:477–522, 2007.

[5] S. Dedeurwaerder, M. Defrance, M. Bizet, and others. A comprehensive overview of Infinium HumanMethylation450 data processing. *Briefings in Bioinformatics*, 15:929–941, 2014.

[6] L. Devroye. *Non-Uniform random variate generation*. Springer, New York, 1986.

[7] Egor Dolzhenko and Andrew D. Smith. Using beta-binomial regression for high-precision differential methylation analysis in multifactor whole-genome bisulfite sequencing experiments. *BMC Bioinformatics*, 15:215, 2014.

[8] Matthias Döring, Gilles Gasparoni, Jasmin Gries, Karl Nordström, Pavlo Lutsik, Jörn Walter, and Nico Pfeifer. Identification and analysis of methylation call differences between bisulfite microarray and bisulfite sequencing data with statistical learning techniques. *BMC Bioinformatics*, 16(Suppl 3):A7, 2015.

[9] P. Du, X. Zhang, C.-C. Huang, and others. Comparison of Beta-value and M-value methods for quantifying methylation levels by microarray analysis. *BMC Bioinformatics*, 11:587, 2010.

[10] S. L. P. Ferrari and F. Cribari-Neto. Beta regression for modelling rates and proportions. *Journal of Applied Statistics*, 31:799–815, 2004.

[11] B. Hofner, A. Mayr, N. Robinzonov, and M. Schmid. Model-based boosting in R: A hands-on tutorial using the R package mboost. *Computational Statistics*, 29:3–35, 2014.

[12] B. Hofner, A. Mayr, and M. Schmid. gamboostLSS: An R package for model building and variable selection in the GAMLSS framework. *Journal of Statistical Software*, 2016. To appear.

[13] T. Hothorn. Boosting – an unusual yet attractive optimiser. *Methods of Information in Medicine*, 53:417–418, 2014.

[14] W. F. Kibble. A two-variate gamma type distribution. *Sankhya*, 5:137–150, 1941.

[15] P. W. Laird. Principles and challenges of genome-wide DNA methylation analysis. *Nature Reviews Genetics*, 11:191–203, 2011.

[16] K. V. Mardia. *Families of Bivariate Distributions*. Griffin, London, 1970.

[17] A. Mayr, N. Fenske, [...], and M. Schmid. Generalized additive models for location, scale and shape for high dimensional data. *Journal of the Royal Statistical Society, Series C*, 61:403–427, 2012.

[18] S. Nadarajah and S. Kotz. Jensen's bivariate gamma distribution: Ratios of components. *Journal of Statistical Computation and Simulation*, 77:349–358, 2007.

[19] Y. Park, M. E. Figueroa, L. S. Rozek, and M. A. Sartor. MethylSig: A whole genome DNA methylation analysis pipeline. *Bioinformatics*, 30:2414–2422, 2014.

[20] A. Portela and M. Esteller. Epigenetic modifications and human disease. *Nature Biotechnology*, 28:1057–1068, 2010.

[21] A. P. Prudnikov, Y. A. Brychkov, and O. I. Marichev. *Integrals and Series*. Gordon and Breach Science Publishers, Amsterdam, 1986.

[22] R. Rigby and D. M. Stasinopoulos. Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society, Series C*, 54:507–554, 2005.

[23] M. Saadati and A. Benner. Statistical challenges of high-dimensional methylation data. *Statistics in Medicine*, 33:5347–5357, 2014.

[24] F. H. Sarkar. *Epigenetics and Cancer*. Springer, Dordrecht, 2013.

[25] M. Schmid, S. Potapov, A. Pfahlberg, and T. Hothorn. Estimation and regularization techniques for regression models with multidimensional prediction functions. *Statistics and Computing*, 20: 139–150, 2010.

[26] M. Schmid, F. Wickler, K. O. Maloney, and others. Boosted beta regression. *PLoS ONE*, 8(4):e61623, 2013.

[27] D. S. Shames, J. D. Minna, and A. F. Gazdar. DNA methylation in health, disease, and cancer. *Current Molecular Medicine*, 7:85–102, 2007.

[28] P. Singmann, D. Shem-Tov, S. Wahl, and others. Characterization of whole-genome autosomal differences of DNA methylation between men and women. *Epigenetics & Chromatin*, 8:43, 2015.

[29] T. J. Triche, D. J. Weisenberger, D. Van Den Berg, and others. Low-level processing of Illumina Infinium DNA Methylation BeadArrays. *Nucleic Acids Research*, 41(7):e90:191–203, 2013.

[30] K. Varley, J. Gertz, K. Bowling, and others. Dynamic DNA methylation across diverse human cell lines and tissues. *Genome Research*, 23:555–567, 2013.

[31] S. Wahl, N. Fenske, [...], and M. Schmid. On the potential of models for location and scale for genome-wide DNA methylation data. *BMC Bioinformatics*, 15:232, 2014.

[32] N. Zhang, H.-J. Wu, W. Zhang, and others. Predicting tumor purity from methylation microarray data. *Bioinformatics*, 31:3401–3405, 2015.

[33] Hao Zheng, Hongwei Wu, Jinping Li, and Shi-Wen Jiang. CpGIMethPred: Computational model for predicting methylation status of CpG islands in human genome. *BMC Medical Genomics*, 66(Suppl 1):S13, 2013.

Appendix

Proof of Proposition 1

We start with a lemma on the properties of the modified Bessel function of the first kind of order $\nu := \alpha - 1$.

Lemma 1: For $\tilde{\alpha} + \nu > 0$ and $p > c$ it holds that

$$\int_0^\infty x^{\tilde{\alpha}-1} \exp(-px)\, I_\nu(cx)\, dx \;=\; p^{-(\tilde{\alpha}+\nu)} \left(\frac{c}{2}\right)^\nu \frac{\Gamma(\nu+\tilde{\alpha})}{\Gamma(\nu+1)}\; {}_2F_1\left(\frac{\nu+\tilde{\alpha}}{2}, \frac{\nu+\tilde{\alpha}+1}{2}, \nu+1, \frac{c^2}{p^2}\right),$$

where ${}_2F_1(\cdot)$ is the Gauss hypergeometric function (see [18], p. 350). For a formal proof of Lemma 1, see [21]. □

The proof of Proposition 1 is obtained by deriving the joint density function $f_{R,b}$ of the random variables $R := M + U$ and $b = M/(M+U) = M/R$. Transforming $(M, U) = (Rb, R(1-b))$ into (R, b) yields the Jacobian matrix

$$\tilde{J} = \begin{pmatrix} \frac{\partial Rb}{\partial R} & \frac{\partial Rb}{\partial b} \\ \frac{\partial R(1-b)}{\partial R} & \frac{\partial R(1-b)}{\partial b} \end{pmatrix} = \begin{pmatrix} b & R \\ (1-b) & -R \end{pmatrix} \tag{15}$$

with $|\det(\tilde{J})| = R$. It follows that, under the assumptions of Proposition 1,

$$\begin{aligned} f_{R,b}(r,b) \;&=\; \frac{(\lambda_m\lambda_u)^{\frac{\alpha+1}{2}}}{(1-\rho)\,\rho^{\frac{\alpha-1}{2}}\Gamma(\alpha)}\; r^\alpha\, (b\,(1-b))^{\frac{\alpha-1}{2}} \exp\left(-\frac{\lambda_m r b + \lambda_u r(1-b)}{1-\rho}\right) \\ &\quad \times I_{\alpha-1}\left(\frac{2\sqrt{\rho\,\lambda_m\lambda_u\, r^2\, b\,(1-b)}}{1-\rho}\right). \end{aligned} \tag{16}$$

Defining

$$Z(b) := \int r^\alpha \exp\left(-\frac{\lambda_m r b + \lambda_u r(1-b)}{1-\rho}\right) I_{\alpha-1}\left(\frac{2\sqrt{\rho\,\lambda_m\lambda_u\, r^2\, b\,(1-b)}}{1-\rho}\right) dr\,, \tag{17}$$

the marginal density function $f_b(b)$ is derived by integrating $f_{R,b}$ over R:

$$f_b(b) = \int f_{R,b}(r,b)\, dr = \frac{(\lambda_m\lambda_u)^{\frac{\alpha+1}{2}}}{(1-\rho)\,\rho^{\frac{\alpha-1}{2}}\,\Gamma(\alpha)}\, (b\,(1-b))^{\frac{\alpha-1}{2}}\, Z(b)\,. \tag{18}$$

Setting

$$\tilde{\alpha} = \alpha + 1\,, \quad \nu = \alpha - 1\,, \quad p = \frac{\lambda_m b + \lambda_u(1-b)}{1-\rho}\,, \quad c = \frac{2\sqrt{\rho\,\lambda_m\lambda_u\, b\,(1-b)}}{1-\rho} \tag{19}$$

and making use of the fact that

$${}_2F_1\,(\alpha, \delta, \alpha, x) = (1-x)^{-\delta}\,, \tag{20}$$

one obtains by application of Lemma 1 that

$$Z(b) = \frac{\Gamma(2\alpha)}{\Gamma(\alpha)}\,(1-\rho)^{\alpha+1}\, \frac{\left(\sqrt{\rho\,\lambda_m\lambda_u\, b\,(1-b)}\right)^{\alpha-1}}{(\lambda_m b + \lambda_u(1-b))^{2\alpha}} \left(1 - \frac{4\rho\,\lambda_m\lambda_u\, b\,(1-b)}{(\lambda_m b + \lambda_u(1-b))^2}\right)^{-\frac{2\alpha+1}{2}}. \tag{21}$$

Combining (18) and (21) yields the probability density function stated in Proposition 1. □

Proof of Proposition 2

Defining $\theta := \lambda_m/\lambda_u$, the log-likelihood function derived from of Equation (9) of the manuscript becomes

$$\begin{aligned}
\sum_{i=1}^{n} \log(f_b(b_i;\alpha,\rho,\theta)) &= \sum_{i=1}^{n} \Big[\log(\Gamma(2\alpha)) - 2\log(\Gamma(\alpha)) + \alpha\log(\lambda_m\lambda_u) + \alpha\log(1-\rho) \\
&\quad + (\alpha-1)\log(b_i(1-b_i)) + \log(\lambda_m b_i + \lambda_u(1-b_i)) \\
&\quad - (\alpha+0.5)\log\left((\lambda_m b_i + \lambda_u(1-b_i))^2 - 4\,\rho\,\lambda_m\lambda_u\,b_i(1-b_i)\right)\Big] \\
&= \sum_{i=1}^{n} \Big[\log(\Gamma(2\alpha)) - 2\log(\Gamma(\alpha)) + \alpha\log(\theta\lambda_u^2) + \alpha\log(1-\rho) \\
&\quad + (\alpha-1)\log(b_i(1-b_i)) + \log\left(\lambda_u\left((\theta-1)\,b_i+1\right)\right) \\
&\quad - (\alpha+0.5)\log\left(\left(\lambda_u((\theta-1)\,b_i+1)\right)^2 - 4\,\rho\,\theta\lambda_u^2\,b_i(1-b_i)\right)\Big] \\
&= \sum_{i=1}^{n} \Big[\log(\Gamma(2\alpha)) - 2\log(\Gamma(\alpha)) + \alpha\log(\theta) + \alpha\log(1-\rho) \\
&\quad + (\alpha-1)\log(b_i(1-b_i)) + \log\left((\theta-1)\,b_i+1\right) \\
&\quad - (\alpha+0.5)\log\left(((\theta-1)\,b_i+1)^2 - 4\,\rho\,\theta\,b_i(1-b_i)\right)\Big]. \qquad (22)
\end{aligned}$$

□

Derivation of the Observed Information Matrix

Defining $D_1 := (\exp(\boldsymbol{X}_i^T\gamma)-1)\,b_i + 1$ and $D_2 := D_1^2 - 4\,\rho\exp(\boldsymbol{X}_i^T\gamma)\,b_i\,(1-b_i)$, the first derivative of Equation (13) of the manuscript w.r.t. γ is given by

$$\begin{aligned}
&\frac{\partial}{\partial\gamma}\sum_{i=1}^{n}\log(f_b(b_i,\boldsymbol{X}_i;\alpha,\rho,\gamma)) \\
&= \sum_{i=1}^{n}\Big[\alpha\boldsymbol{X}_i^T + \frac{\boldsymbol{X}_i^T b_i\exp(\boldsymbol{X}_i^T\gamma)}{D_1} - (\alpha+0.5)\,\frac{2\,(D_1 - 2\,\rho\,(1-b_i))\,\boldsymbol{X}_i^T b_i\exp(\boldsymbol{X}_i^T\gamma)}{D_2}\Big]. \qquad (23)
\end{aligned}$$

It follows that the observed information matrix is given by

$$\begin{aligned}
J(\alpha,\rho,\gamma) &= -\frac{\partial^2}{\partial^2\gamma}\sum_{i=1}^{n}\log(f_b(b_i,\boldsymbol{X}_i;\alpha,\rho,\gamma)) \qquad (24) \\
&= \sum_{i=1}^{n}\Big[\frac{\left(D_1 - b_i\exp(\boldsymbol{X}_i^T\gamma)\right)b_i\exp(\boldsymbol{X}_i^T\gamma)\boldsymbol{X}_i\boldsymbol{X}_i^T}{D_1^2} \\
&\quad - (\alpha+0.5)\,\frac{2\,b_i\exp(\boldsymbol{X}_i^T\gamma)\left(D_1 + b_i\exp(\boldsymbol{X}_i^T\gamma) - 2\,\rho\,(1-b_i)\right)\boldsymbol{X}_i\boldsymbol{X}_i^T}{D_2} \\
&\quad + (\alpha+0.5)\,\frac{4\,b_i^2\exp(2\boldsymbol{X}_i^\top\gamma)\,(D_1 - 2\rho(1-b_i))^2\,\boldsymbol{X}_i\boldsymbol{X}_i^T}{D_2^2}\Big]. \qquad (25)
\end{aligned}$$